arXiv:1601.01975v1 [quant-ph] 8 Jan 2016

Quantum Merlin Arthur with Exponentially Small Gap

Bill Fefferman[1] and Cedric Lin[1]

[1]Joint Center for Quantum Information and Computer Science (QuICS), University of Maryland

January 11, 2016

Abstract

We will study the complexity of QMA proof systems with inverse exponentially small promise gap. We will show that this class, $\mathsf{QMA}_{\mathsf{exp}}$, can be exactly characterized by PSPACE, the class of problems solvable with a polynomial amount of memory. As applications we show that a "precise" version of the Local Hamiltonian problem is PSPACE-complete, and give a provable setting in which the ability to prepare PEPS states is not as powerful as the ability to prepare the ground state of general Local Hamiltonians.

1 Introduction

The class QMA, *Quantum Merlin-Arthur*, is the quantum analogue of NP, and is one of the central objects of study in quantum complexity theory. QMA consists of those problems whose solutions can be verified with high probability using a quantum computer. This class was first shown to have a natural complete problem, the Local Hamiltonian problem, in [10]; since then many more QMA-complete problems have been discovered (see e.g., [7]). There has also been much work in trying to prove a quantum version of the PCP theorem; see [2] for a review.

To be more precise, we give here the definition of QMA:

Definition 1. *We say a promise problem $L = (L_{yes}, L_{no})$ is in $\mathsf{QMA}(c,s)$ if there exists a uniform family of quantum circuits $\{V_x\}_{x\in\{0,1\}^n}$, each of at most polynomial size, and acting on $k(|x|)+m(|x|)$ qubits for polynomials k and m, so that:*

If $x \in L_{yes}$ there exists an m-qubit state $|\psi\rangle$ such that:

$$\left(\langle\psi| \otimes \left\langle 0^k\right|\right) V_x^\dagger |1\rangle\langle 1|_{out} V_x \left(|\psi\rangle \otimes \left|0^k\right\rangle\right) \geq c \tag{1}$$

Whereas if $x \in L_{no}$, for all m-qubit states $|\psi\rangle$ we have:

$$\left(\langle\psi| \otimes \left\langle 0^k\right|\right) V_x^\dagger |1\rangle\langle 1|_{out} V_x \left(|\psi\rangle \otimes \left|0^k\right\rangle\right) \leq s. \tag{2}$$

We call c the completeness and s the soundness parameters. Then $\mathsf{QMA} = \mathsf{QMA}(2/3, 1/3)$.

It is natural to wonder whether the precise values of c and s matter. Kitaev showed [10] that as long as c and s are separated by at least an inverse polynomial, then by repeating the verification circuit polynomially many times, it is possible to amplify the promise gap $c-s$ to any constant less than one. Thus $\mathsf{QMA}(c, c-1/\mathrm{poly}) = \mathsf{QMA}$, and quantum Merlin Arthur proof protocols with only a polynomial gap is just as powerful as QMA.

1.1 Our contribution

In this work we study the complexity of QMA protocols where the gap is only exponentially small, i.e. $c - s = \exp(-\text{poly})$. We show that in this case, the problems verifiable by these protocols exactly coincide with the problems solvable in classical polynomial space:

Theorem 1. $\mathsf{QMA}_{\mathsf{exp}} := \cup_{c-s>\exp(-\text{poly})}\mathsf{QMA}(c,s) = \mathsf{PSPACE}$.

The closest classical counterpart of $\mathsf{QMA}_{\mathsf{exp}}$ is $\mathsf{NP}^{\mathsf{PP}}$: given a classical witness, the verifier runs a classical computation that in the YES case accepts with probability at least c, or in the NO case accepts with probability at most s, where $c > s$. Note that in the classical case the inequality $c - s > \exp(-\text{poly})$ is always satisfied. Since $\mathsf{NP}^{\mathsf{PP}}$ is in the counting hierarchy, the entirety of which is contained in PSPACE (see e.g., [3]), we see that the quantum proof protocol is strictly stronger than the classical one, unless the counting hierarchy collapses to the second level.

Our proof of this theorem allows us to tweak the proof of QMA-completeness of Local Hamiltonian [10, 9] to show the following:

Theorem 2. *For any $3 \leq k \leq \mathcal{O}(\log(n))$, determining whether the ground state energy of a k-local Hamiltonian is at most a or at least b for $b - a > \exp(-\text{poly})$ is* PSPACE*-complete.*[1]

In contrast, when $b - a > 1/\text{poly}$ the k-Local Hamiltonian problem is QMA-complete [10, 9] for $2 \leq k \leq \mathcal{O}(\log(n))$.

Moreover, since we can binary search for the ground state energy in classical polynomial space, this shows

Corollary 3. *For any $3 \leq k \leq \mathcal{O}(\log(n))$, computing the ground state energy of a k-local Hamiltonian to polynomially many bits of precision is a* FPSPACE*-complete problem.*

Here recall that FPSPACE is the set of functions computable in classical polynomial space.

We have therefore shown that there is, unsurprisingly, a large jump in complexity for the local Hamiltonian problem when the promise gap is only exponentially small instead of polynomially small. Perhaps more surprisingly, $\mathsf{QMA}_{\mathsf{exp}} = \mathsf{PSPACE}$ is more powerful than $\mathsf{PostBQP} = \mathsf{PP}$, the class of problems solvable with postselected quantum computation [1].

Recall that projected entangled pair states, or PEPS, are a natural extension of matrix product states to two and higher dimensions, and can described as the ground state of certain frustration-free local Hamiltonians [15]. A characterization of the computational power of PEPS was given in [14], and can be summarized as follows: let O_{PEPS} be a quantum oracle that, given the description of a PEPS, outputs the PEPS (so the output is quantum). Then $\mathsf{BQP}^{O_{PEPS}}_{\|,\text{classical}} = \mathsf{PostBQP} = \mathsf{PP}$, where (following Aaronson [1]) the subscript denotes that only classical nonadaptive queries to the oracle are allowed. Moreover, let PQP stand for the set of problems solvable by a quantum computer with *unbounded error*; then it can be straightforwardly shown that $\mathsf{PQP}^{O_{PEPS}}_{\|,\text{classical}} = \mathsf{PP}$ as well (see Appendix A for a proof sketch).

On the other hand, suppose we have an oracle O_{LH} that given the description of a local Hamiltonian, outputs its unique ground state[2]. Then our results show that $\mathsf{QMA}_{\mathsf{exp}} = \mathsf{PSPACE} \subseteq \mathsf{PQP}^{O_{LH}}_{\|,\text{classical}}$. This shows that, at least in the setting of unbounded-error quantum computation, PEPS do not capture the full computational complexity of general local Hamiltonian ground states unless $\mathsf{PP} = \mathsf{PSPACE}$. We leave open the problem of determining the complexity of $\mathsf{BQP}^{O_{LH}}_{\|,\text{classical}}$.

[1] We make no attempt to decide whether PSPACE-completeness still holds for $k = 2$.

[2] If the ground space is degenerate, we can always slightly perturb the Hamiltonian to make it nondegenerate.

2 Definitions

2.1 Quantum Merlin Arthur

For our purposes we will need to keep track of the time and space requirements of QMA protocols, and so we make the following definition:

Definition 2. *We say a promise problem $L = (L_{yes}, L_{no})$ is in (t,k)-bounded $\mathsf{QMA}_m(c,s)$ if there exists a uniform family of quantum circuits $\{V_x\}_{x\in\{0,1\}^n}$, each of size at most $t(|x|)$, acting on $k(|x|)+m(|x|)$ qubits, so that:*

If $x \in L_{yes}$ there exists an m-qubit state $|\psi\rangle$ such that:

$$\left(\langle\psi| \otimes \left\langle 0^k\right|\right) V_x^\dagger |1\rangle\langle 1|_{out} V_x \left(|\psi\rangle \otimes \left|0^k\right\rangle\right) \geq c \tag{3}$$

Whereas if $x \in L_{no}$, for all m-qubit states $|\psi\rangle$ we have:

$$\left(\langle\psi| \otimes \left\langle 0^k\right|\right) V_x^\dagger |1\rangle\langle 1|_{out} V_x \left(|\psi\rangle \otimes \left|0^k\right\rangle\right) \leq s. \tag{4}$$

Remark 4. $\mathsf{QMA} = (\mathrm{poly}, \mathrm{poly})$-*bounded* $\mathsf{QMA}_{\mathrm{poly}}(2/3, 1/3)$

Remark 5. $\mathsf{QMA}_{\mathsf{exp}} = (\mathrm{poly}, \mathrm{poly})$-*bounded* $\mathsf{QMA}_{\mathrm{poly}}(c, c - 2^{-\mathrm{poly}})$

2.2 Space complexity classes

Definition 3. *We say a promise problem $L = (L_{yes}, L_{no})$ is in* PQPSPACE *(unbounded-error quantum polynomial space) if there exists a uniform family of quantum circuits acting on at most polynomial number of qubits that accepts every string $x \in L_{yes}$ with probability at least c, and accepts every string $x \in L_{no}$ with probability at most s, for some $c > s$.*

Theorem 6 (Watrous [16, 17])**.** PQPSPACE = PSPACE.

Definition 4. *We say a promise problem $L = (L_{yes}, L_{no})$ is in* revPSPACE *(reversible polynomial space) if it can be decided by a polynomial space reversible Turing machine, i.e., a machine for which every configuration has at most one immediate predecessor.*

Theorem 7 (Bennett [4])**.** revPSPACE = PSPACE.

3 Upper bound

In this section, our goal will be to prove that $\mathsf{QMA}_{\mathsf{exp}} \subseteq \mathsf{PSPACE}$. We will proceed in two steps, the first will show how to use in-place QMA amplification techniques from Nagaj, Wocjan, and Zhang [13] to decide any promise problem in $\mathsf{QMA}_{\mathsf{exp}}$ with a quantum protocol in which the verifier is allowed exponential time, polynomial space (i.e., acts on a polynomial number of proof and ancilla qubits), completeness $1 - 2^{-\mathrm{poly}(n)}$ and soundness $2^{-\mathrm{poly}(n)}$. We then appeal to results of Marriott and Watrous [12] and Watrous [16] to show that such protocols can be simulated in PSPACE.

3.1 In-place gap amplification of $\mathsf{QMA}_{\mathsf{exp}}$ using phase estimation techniques

Theorem 8 (Implicit in Nagaj, Wocjan, and Zhang [13])**.** *For any $r > 0$,*

$$(t,k)\text{-bounded } \mathsf{QMA}_m(c,s) \subseteq (\mathcal{O}(rt(c-s)), \mathcal{O}(k + r\log(c-s)))\text{-bounded } \mathsf{QMA}_m(1-2^{-r}, 2^{-r}).$$

Proof. Let $L = (L_{yes}, L_{no})$ be a promise problem in $\mathsf{QMA}(c,s)$ and $\{V_x\}_{x\in\{0,1\}^n}$ the corresponding uniform family of verification circuits. Define the projectors:

$$\Pi_0 = I_m \otimes \left|0^k\right\rangle\!\left\langle 0^k\right| \tag{5}$$

$$\Pi_1 = V_x^\dagger \left(|1\rangle\langle 1|_{out} \otimes I_{m+k-1}\right) V_x \tag{6}$$

and the corresponding reflections:

$$R_0 = 2\Pi_0 - I \tag{7}$$

$$R_1 = 2\Pi_1 - I. \tag{8}$$

Now consider the following procedure:

1. Perform r trials of phase estimation of the operator R_1R_0 on the state $|\psi\rangle \otimes \left|0^k\right\rangle$, with $\mathcal{O}(\log(c-s))$ bits of precision and $1/16$ failure probability.
2. If the median of the r results is at most $\phi_c = \arccos\sqrt{c}/\pi$, output YES; otherwise if the result is at least $\phi_s = \arccos\sqrt{s}/\pi$, output NO.

Phase estimation of an operator U up to a bits of precision requires $\mathcal{O}(a)$ ancilla qubits and $\mathcal{O}(2^a)$ applications of the control-U operation. Thus, the above procedure, which uses r applications of phase estimation to precision $\alpha = \mathcal{O}(\log(c-s))$ on the V_x operator, can be implemented by a circuit of size $\mathcal{O}(rt2^\alpha) = \mathcal{O}(rt(c-s))$ using $\mathcal{O}(r\alpha)$ extra ancillia qubits. Using the standard analysis of in-place QMA error amplification [12, 13], it can be seen that this procedure has completeness probability at least $1-2^{-r}$ and soundness at most 2^{-r}. □

Thus, we get the following corollaries:

Corollary 9. *For all $r > 0$, $\mathsf{QMA}_{\mathsf{exp}} \subseteq \left(r2^{\mathrm{poly}}, r\cdot\mathrm{poly}\right)$-bounded $\mathsf{QMA}_{\mathrm{poly}}(1-2^{-r}, 2^{-r})$.*

Corollary 10. *For every problem $L \in \mathsf{QMA}_{\mathsf{exp}}$, there exists an $m \in \mathrm{poly}$ so that:*

$$L \in \left(2^{\mathrm{poly}}, \mathrm{poly}\right)\text{-bounded } \mathsf{QMA}_m(1-2^{-(m+2)}, 2^{-(m+2)})$$

Notice that Corollary 10 follows from the definition of $\mathsf{QMA}_{\mathsf{exp}}$ and Corollary 9 with $r = m+2$.

3.2 PSPACE simulation

Theorem 11. *For all $m \in \mathrm{poly}$:*

$$\left(2^{\mathrm{poly}}, \mathrm{poly}\right)\text{-bounded } \mathsf{QMA}_m(1-2^{-(m+2)}, 2^{-(m+2)}) \subseteq \mathsf{PQPSPACE}$$

Proof. For any $m, p \in \text{poly}$, consider a problem $L \in \left(2^{\text{poly}}, p\right)$-*bounded* $\mathsf{QMA}_m(1-2^{-(m+2)}, 2^{-(m+2)})$, and let $\{V'_x\}_{x\in\{0,1\}^n}$ be the corresponding uniform family of verification circuits. If $x \in L_{yes}$ there exists an m-qubit state $|\psi\rangle$ such that

$$\left(\langle\psi| \otimes \langle 0^p|\right) V_x'^{\dagger} |1\rangle\langle 1|_{out} V'_x \left(|\psi\rangle \otimes |0^p\rangle\right) \geq 1 - 2^{-(m+2)} \tag{9}$$

whereas if $x \in L_{no}$, for all m-qubit states $|\psi\rangle$ we have

$$\left(\langle\psi| \otimes \langle 0^p|\right) V_x'^{\dagger} |1\rangle\langle 1|_{out} V'_x \left(|\psi\rangle \otimes |0^p\rangle\right) \leq 2^{-(m+2)}. \tag{10}$$

For convenience, define the $2^m \times 2^m$ matrix:

$$Q_x := \left(I_{2^m} \otimes \langle 0^p|\right) V_x'^{\dagger} |1\rangle\langle 1|_{out} V'_x \left(I_{2^m} \otimes |0^p\rangle\right). \tag{11}$$

Q_x is positive semidefinite, and $\langle\psi|Q_x|\psi\rangle$ is the acceptance probability of V'_x on witness ψ. Note that

$$x \in L_{yes} \Rightarrow \text{tr}[Q_x] \geq 1 - 2^{-(m+2)} \geq 3/4 \tag{12}$$

since the trace is at least the largest eigenvalue, and $m \geq 0$; likewise,

$$x \in L_{no} \Rightarrow \text{tr}[Q_x] \leq 2^m \cdot 2^{-(m+2)} = 1/4 \tag{13}$$

since the trace is the sum of the 2^m eigenvalues, each of which is at most $2^{-(m+2)}$.

Therefore our problem reduces to determining whether the trace of Q_x is at least $3/4$ or at most $1/4$. Now we will show that using the totally mixed state $2^{-m} I_m$ (alternatively, a random computational basis state) as the witness of the verification procedure encoded by Q_x, succeeds with the desired completeness and soundness bounds. The acceptance probability is given by

$$\text{tr}(Q_x 2^{-m} I_m) = 2^{-m} \,\text{tr}(Q_x) \tag{14}$$

which is at least $2^{-m} \cdot 3/4$ if $x \in L_{yes}$, and at most $2^{-m} \cdot 1/4$ if $x \in L_{no}$. Thus we have reduced our original problem to determining whether an exponentially long quantum computation with *no* witness, acting on a polynomial number of qubits, accepts with probability at least c' or at most s' with $c' - s'$ being exponentially small. This is a PQPSPACE problem.

□

Theorem 12. $\mathsf{QMA}_{\mathsf{exp}} \subseteq \mathsf{PSPACE}$.

Proof. This follows from Theorem 6, Corollary 10, and Theorem 11. □

4 Lower bound

In this section we will show that $\mathsf{PSPACE} \subseteq \mathsf{QMA}_{\mathsf{exp}}$. To do this we proceed with two steps. In the first we show that, given a succinctly representable sparse matrix and promised that either the smallest eigenvalue is 0 or at most $1/2^{\text{poly}}$, deciding which is the case is a PSPACE-complete problem. In the second step, we give a $\mathsf{QMA}_{\mathsf{exp}}$ protocol for this problem.

By a matrix being succinctly representable and sparse, we mean the following:

Definition 5. *Let M be a $2^{\text{poly}(n)} \times 2^{\text{poly}(n)}$ matrix, where n is the input size. We say that M is a* succinctly representable sparse *matrix if there are at most polynomially many nonzero entries in each row, and moreover there is a (uniformly generated) circuit,* the succinct encoding*, that outputs the nonzero entries of any given row in* $\text{poly}(n)$ *time.*

4.1 The Succinct Determinant problem

Definition 6 (*Succinct Determinant Checking*)**.** *Given as input is a succinct encoding of A, a succinctly representable sparse matrix, whose determinant is promised to be 0, 1, or -1. Moreover, each column of A has at no more than two 1s. Does* $\det(A)$ *vanish?*

Theorem 13 (Grenet, Koiran and Portier [8])**.** *Succinct Determinant Checking is* PSPACE*-hard.*

Proof. Let $L = (L_{yes}, L_{no}) \in$ PSPACE be decided by a polynomial space deterministic Turing Machine M. Consider the configuration graph G^M of M: each vertex of G^M corresponds to one of the exponentially many configurations of M, each of which is describable with polynomially many bits. The configuration graph G^M has an edge from c to c' if and only if c' can be reached from c in one step of computation. It is straightforward to see that G^M has the following properties:

- Since M is deterministic, all vertices of G^M have out-degree at most 1, and G_M has no cycles.
- The adjacency matrix of G^M is a succinctly representable sparse matrix.
- M accepts input x if and only if there is a path in G^M from the starting configuration s_x to the accepting configuration t.

Now on input x, consider the graph G_x^M obtained by adding an edge from the accepting configuration t to the starting configuration s_x, and adding self-loops on all other vertices.

Recall that a *cycle cover* of a directed graph is a set of disjoint cycles that are subgraphs containing all the vertices of the graph. We define the *signed weight* of a cycle cover to be the product of the weights of the edges in the cycle cover, multiplied by $(-1)^\ell$, where ℓ is the number of cycles of even length in the cycle cover. We can interpret the determinant of the adjacency matrix of a directed graph as the sum of the signed weights of all cycle covers of the graph.

Let A_x^M be the adjacency matrix of G_x^M. Now G_x^M has a cycle cover if and only if M accepts x and there is a path from s_x to t, and so $\det(A_x^M) = \pm 1$, depending on the signed weight of the cycle cover; otherwise, $\det(A_x^M) = 0$. Therefore deciding whether $\det(A_x^M)$ vanishes is PSPACE-hard. □

We can immediately see that the complexity of the *Succinct Determinant Checking* problem doesn't get easier if we are promised that the succinctly representable sparse input matrix A is symmetric and positive semidefinite: notice that the matrix $(A_x^M)^T A_x^M$ is succinctly representable sparse because there are at most two 1's in each column of A_x^M and if we can decide if $\det((A_x^M)^T A_x^M) = \det(A_x^M)^2$ vanishes (or is equal to 1), we can certainly decide if $\det(A_x^M)$ vanishes.

From here, we would like to argue that given a succinctly representable sparse and symmetric PSD matrix A, it is PSPACE-hard to determine whether the smallest eigenvalue λ_{min} satisfies $\lambda_{min} = 0$ or $\lambda_{min} > 2^{-\text{poly}}$, promised that one of these is the case. We will see later that this problem can be solved in $\mathsf{QMA}_{\mathsf{exp}}$. Unfortunately, this promise does not generally hold for succinctly representable sparse symmetric matrices; if A is nonsingular, the smallest eigenvalue can at worst still be doubly exponentially small. We will therefore need to modify the prior PSPACE-hard construction to show that the following *Gapped Succinct Matrix Singularity* problem is still PSPACE-hard.

Definition 7 (*Gapped Succinct Matrix Singularity*)**.** *Given as input is a succinct encoding of A, a positive semidefinite, symmetric, and succinctly representable sparse matrix, whose entries are 0, 1, or 2. Moreover, the smallest eigenvalue of A is promised to be either zero or at least $2^{-g(n)}$ for some polynomial $g(n)$. Output YES if the smallest eigenvalue of A is zero.*

Theorem 14. *Gapped Succinct Matrix Singularity is* PSPACE*-hard.*

Proof. In this theorem, we will adapt the construction of Theorem 13 to analyze the spectrum of the underlying PSPACE machine configuration graph. Recall we defined revPSPACE to be the class of languages decidable in polynomial space by a reversible Turing Machine. Theorem 7 states that PSPACE = revPSPACE, and indeed a result of Lange, McKenzie and Tapp [11] proves that for any space-constructible s, DSPACE[s] $\subseteq$ revSPACE[s], at a cost of an exponential time blow-up. In fact, they also show without loss of generality that the starting configuration of the reversible machine has in-degree 0, as long as the input x is kept on the tape at the end of the computation (and so the accepting configuration depends on x).

Thus, an arbitrary PSPACE language L can be decided by a polynomial space reversible Turing machine M, and the resulting configuration graph G^M is a collection of disjoint paths.

As before consider the graph G_x^M obtained by adding an edge from the accepting configuration t to the starting configuration s_x, and adding self-loops on all other vertices. Now, if $\det(A_x^M) \neq 0$, i.e. if M accepts x, there is a maximal path in G^M starting from s_x and ending at t_x. Assume the path has $\ell + 1$ vertices, where $\ell \in 2^{\text{poly}}$. G_x^M adds an edge from t_x to s_x and adds a self-loop to all other vertices. Therefore if M accepts x, G_x^M is a disjoint union of connected graphs $G_{x,i}^M$, where:

1. If s_x, t_x are not vertices of $G_{x,i}^M$, $G_{x,i}^M$ is a path with additional self-loops on all vertices of the path.

2. If s_x, t_x are vertices of $G_{x,i}^M$, $G_{x,i}^M$ is a cycle, with s_x coming directly after t_x in the cycle, and with additional self-loops on all vertices in the cycle except for s_x and t_x.

Let us look at these two cases separately. Assume a subgraph of the first type has ℓ vertices (i.e. the path has length $\ell - 1$); then its adjacency matrix is, after appropriate relabelling of vertices, the following $\ell \times \ell$ matrix:

$$A_{1,\ell} := \begin{bmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 \\ 1 & 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & 0 & \dots & 1 & 1 \end{bmatrix}. \tag{15}$$

Computing $A_{1,\ell}^T A_{1,\ell}$, we see that it is:

$$A_{1,\ell}^T A_{1,\ell} := \begin{bmatrix} 2 & 1 & 0 & 0 & \dots & 0 & 0 \\ 1 & 2 & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 2 & 1 & \dots & 0 & 0 \\ 0 & 0 & 1 & 2 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 2 & 1 \\ 0 & 0 & 0 & 0 & \dots & 1 & 1 \end{bmatrix}. \tag{16}$$

The characteristic equation for the eigenvalues λ is then $p_\ell(\lambda) = 0$, where

$$p_\ell(\lambda) := \det\left(\begin{bmatrix} 2-\lambda & 1 & 0 & 0 & \dots & 0 & 0 \\ 1 & 2-\lambda & 1 & 0 & \dots & 0 & 0 \\ 0 & 1 & 2-\lambda & 1 & \dots & 0 & 0 \\ 0 & 0 & 1 & 2-\lambda & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \\ 0 & 0 & 0 & 0 & \dots & 2-\lambda & 1 \\ 0 & 0 & 0 & 0 & \dots & 1 & 1-\lambda \end{bmatrix}\right). \tag{17}$$

Now define the polynomial $q_n(x)$ to be the following $n \times n$ determinant:

$$q_n(x) := \det\left(\begin{bmatrix} x & 1 & 0 & \dots & 0 & 0 \\ 1 & x & 1 & \dots & 0 & 0 \\ 0 & 1 & x & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \\ 0 & 0 & 0 & \dots & x & 1 \\ 0 & 0 & 0 & \dots & 1 & x \end{bmatrix}\right) \tag{18}$$

Note that $p_\ell(\lambda) = q_\ell(2-\lambda) - q_{\ell-1}(2-\lambda)$. Moreover, $q_n(x)$ satisfies the recurrence

$$q_0(x) = 1, \quad q_1(x) = x, \quad q_n(x) = xq_{n-1}(x) - q_{n-2}(x) \tag{19}$$

This is the same recurrence satisfied by $U_n(x/2)$, where $U_n(x)$ is the n-th degree Chebyshev polynomial of the second kind. Therefore $q_n(x) = U_n(x/2)$. Using $U_n(\cos\theta) = \sin((n+1)\theta)/\sin(\theta)$ we can evaluate:

$$q_\ell(2\cos\theta) - q_{\ell-1}(2\cos\theta) = \frac{\sin((\ell+1)\theta) - \sin(\ell\theta)}{\sin\theta} \tag{20}$$

$$= \frac{\sin((2\ell+1)\theta/2)}{\sin(\theta/2)} \tag{21}$$

and therefore the zeroes of $q_\ell(x) - q_{\ell-1}(x)$ are $2\cos\left(\frac{2k}{2\ell+1}\pi\right)$, $k = 1, \cdots, \ell$. The zeroes of the polynomial $p_\ell(\lambda) = q_\ell(2-\lambda) - q_{\ell-1}(2-\lambda)$ are then

$$\lambda_k = 2\left(1 - \cos\left(\frac{2k}{2\ell+1}\pi\right)\right) \tag{22}$$

and the smallest eigenvalue $\lambda_1 = \Theta(\ell^{-2})$ is inverse exponentially bounded away from zero, because $\ell = 2^{\mathcal{O}(\text{poly})}$.

We now look at the other case, where the subgraph is of the second type, i.e. it contains s_x and t_x. The adjacency matrix of this subgraph is, assuming there are ℓ vertices, the $\ell \times \ell$ matrix:

$$A_{2,\ell} := \begin{bmatrix} 0 & 0 & 0 & 0 & \dots & 0 & 1 \\ 1 & 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & 1 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 & 0 \\ 0 & 0 & 0 & 0 & \dots & 1 & 0 \end{bmatrix} \tag{23}$$

We can directly evaluate $A_{2,\ell}^T A_{2,\ell}$, obtaining the following matrix:

$$A_{2,\ell}^T A_{2,\ell} = \begin{bmatrix} 1 & 1 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1 & 2 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1 & 2 & 1 & \dots & 0 & 0 & 0 \\ 0 & 0 & 1 & 2 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \\ 0 & 0 & 0 & 0 & \dots & 2 & 1 & 0 \\ 0 & 0 & 0 & 0 & \dots & 1 & 2 & 0 \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 1 \end{bmatrix} \tag{24}$$

For purposes of calculating the smallest eigenvalue the last row and column can be ignored, leaving an $(\ell - 1) \times (\ell - 1)$ matrix. Looking back at 16 and 17, we see that the characteristic equation is exactly $p_{\ell-1}(\lambda) = 0$. Therefore once again the smallest eigenvalue is inverse exponentially bounded away from zero.

□

4.2 *Gapped Succinct Matrix Singularity* is in $\mathsf{QMA}_{\mathsf{exp}}$

In this section, we now proceed to show a $\mathsf{QMA}_{\mathsf{exp}}$ protocol for *Gapped Succinct Matrix Singularity*.

Lemma 1. *Let A be a positive semidefinite, symmetric, and succinctly representable sparse matrix, whose entries are 0, 1, or 2; moreover the smallest eigenvalue of A is promised to be either zero or at least $2^{-g(n)}$ for some polynomial $g(n)$. There is a $\mathsf{QMA}_{\mathsf{exp}}$ protocol for deciding which is the case.*

Our strategy will essentially be to simulate the time evolution of the sparse Hamiltonian e^{-iAt} using known simulation methods, and then use a stripped-down version of phase estimation to estimate an eigenvalue of A. We first note the following result for sparse matrix simulation:

Theorem 15 ([5], [6])**.** *Suppose A is a $2^n \times 2^n$ symmetric and succinctly representable sparse matrix, with at most d nonzero entries in each row. Then treated as a Hamiltonian, the time evolution $\exp(-iAt)$ can be simulated using $\mathrm{poly}(n, d, \|A\|, t, \log(1/\epsilon))$ operations.*

The crucial thing to notice in Theorem 15 is the polylogarithmic scaling in the error ϵ; this implies that we can obtain exponential precision in $\exp(-iAt)$ using only polynomially many operations. Also note that we can upper bound $\|A\|$ with the following observation:

Remark 16. *Suppose a matrix A has at most d nonzero entries per row, each of which is no more than k in absolute value. Then $\|A\| \leq kd$.*

Proof of Lemma 1. We are given a succinct encoding of an symmetric PSD d-sparse matrix A, and it is promised that the smallest eigenvalue λ_{min} of A is either zero or at least $2^{-g(n)}$ for some polynomial $g(n)$. Merlin would like to convince us that $\lambda_{min} = 0$; he will send us a purported eigenstate $|\psi\rangle$ of A with zero eigenvalue. We will carry out a stripped-down version of phase estimation on $\exp(-iAt)$ acting on $|\psi\rangle$ to decide, with exponentially small completeness-soundness gap, whether Merlin is telling the truth. Let us choose $t = \pi/(kd) \leq \pi/\|A\|$; then all eigenvalues of At lie in the range $[0, \pi]$, and the output of phase estimation will be unambiguous.

To implement phase estimation, we first need to be able to implement $\exp(-iAt)$ efficiently and to high precision. This is what Theorem 15 gives us: we can implement $\exp(-iAt)$ up to error $\epsilon = 2^{-\text{poly}(n)}$ using only a polynomial number of operations, for any choice of ϵ.

Now to use phase estimation to distinguish the phase up to exponential precision, we would normally require exponentially many operations in the usual phase estimation routine. Instead, we will simply do phase estimation with one bit:

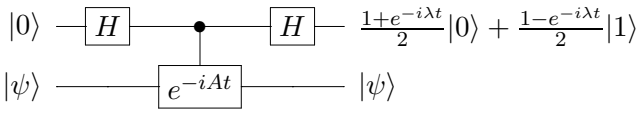

(25)

In the above we've assumed $|\psi\rangle$ is an eigenstate of A with eigenvalue λ. If we measure the control qubit at the end, we see the probability we obtain 0 is $1-(1-\cos(\lambda t))/2 = 1-(\lambda t)^2/4+\mathcal{O}(\lambda^4 t^4)$. Therefore if ψ is a zero eigenstate, we can verify this with probability at least $1-\epsilon$, where recall ϵ is the error in the implementation of $\exp(-iAt)$. Otherwise if $\lambda_{min} \geq 2^{-g(n)}$, no state ψ will be accepted with probability greater than $1-2^{-2g(n)}t^2/4+\epsilon+\mathcal{O}(2^{-4g(n)}t^4)$. The separation between the completeness and soundness probabilities is exponentially small if we pick $\epsilon \leq 2^{-2g(n)}t^2/16$, and this therefore gives us a $\mathsf{QMA}_{\mathsf{exp}}$ protocol. □

Theorem 17. $\mathsf{PSPACE} \subseteq \mathsf{QMA}_{\mathsf{exp}}$.

Proof. Follows from Theorem 14 and Lemma 1. □

This finishes the proof of our main theorem:

Theorem 1. $\mathsf{QMA}_{\mathsf{exp}} := \cup_{c-s>\exp(-\text{poly})}\mathsf{QMA}(c,s) = \mathsf{PSPACE}$.

5 A PSPACE-complete variant of the Local Hamiltonian problem

The classic QMA-complete problem is the Local Hamiltonian problem: given a local Hamiltonian H, and parameters $a < b$ with $b - a > 1/\text{poly}$, it is promised that the smallest eigenvalue of H is either at most a or at least b; decide which is the case. We now show that if we weaken the promise gap from polynomially small to only exponentially small, then this problem becomes PSPACE-complete.

Definition 8 (*Precise k-Local Hamiltonian*)**.** *Given as input is a k-local Hamiltonian $H = \sum_{j=1}^{r} H_j$ acting on n qubits, satisfying $r \in \text{poly}(n)$ and $\|H_j\| \leq \text{poly}(n)$, and numbers $a < b$ satisfying $b - a > 2^{-\text{poly}(n)}$. It is promised that the smallest eigenvalue of H is either at most a or at least b. Output 1 if the smallest eigenvalue of H is at most a, and output 0 otherwise.*

Theorem 2. *For any $3 \leq k \leq \mathcal{O}(\log(n))$, Precise k-Local Hamiltonian is $\mathsf{QMA}_{\mathsf{exp}}$-complete, and hence PSPACE-complete.*

Proof. This proof follows straightforwardly by adapting the proof of [10] and [9]. The proof of containment in $\mathsf{QMA}_{\mathsf{exp}}$ is identical to the containment of the usual Local Hamiltonian problem in QMA; see [10] for details.

To show $\mathsf{QMA}_{\mathsf{exp}}$-hardness, we note that for a QMA-verification procedure with T gates, completeness c and soundness s, [9] reduces this to a 3-local Hamiltonian with lowest eigenvalue no more than $(1-c)/(T+1)$ in the YES case, or no less than $(1-s)/T^3$ in the NO case. For this to specify a valid *Precise Local Hamiltonian* problem we need that

$$\frac{1-s}{T^3} - \frac{1-c}{T+1} > 2^{-\text{poly}(n)}. \tag{26}$$

Fortunately, there are indeed values of c and s that satisfy the above inequality and can still specify $\mathsf{QMA}_{\mathsf{exp}}$-hard problems. To see this, we recall the proof of Lemma 1: there it was shown that any problem in PSPACE can be reduced to a $\mathsf{QMA}_{\mathsf{exp}}$ problem with soundness and completeness

$$1 - c = \epsilon, \quad 1 - s = -\epsilon + 2^{-g'(n)} \tag{27}$$

for some polynomial $g'(n)$ depending on the problem, and any $\epsilon = 2^{-\mathrm{poly}(n)}$ of our choice. The number of operations for that protocol is upper bounded by $T \leq h(n, \log(1/\epsilon))$ for some polynomial $h(x, y)$. Now we can pick ϵ to be a small enough inverse exponential function such that

$$\epsilon(T^2 + 1) \leq \epsilon[h^2(n, \log(1/\epsilon)) + 1] < 2^{-g'(n)} \tag{28}$$

holds; this then implies the inequality 26. Hence any problem in PSPACE can be reduced to a *Precise 3-Local Hamiltonian* problem. □

6 Acknowledgements

We are grateful to Sevag Gharibian and Martin Schwarz for helpful conversations, and to John Watrous for comments on a preliminary draft. This work was supported by the Department of Defense.

References

[1] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A*, 461(2063):3473–3482, 2005.

[2] D. Aharonov, I. Arad, and T. Vidick. The Quantum PCP Conjecture. *ArXiv e-prints*, September 2013.

[3] Eric W Allender and Klaus W Wagner. *Counting hierarchies: polynomial time and constant depth circuits*. World Scientific.

[4] Charles H. Bennett. Time/space trade-offs for reversible computation. *SIAM J. Comput.*, 18(4):766–776, August 1989.

[5] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Exponential improvement in precision for simulating sparse hamiltonians. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 283–292, New York, NY, USA, 2014. ACM.

[6] Dominic W. Berry, Andrew M. Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. 2015.

[7] Adam D. Bookatz. QMA-complete problems. *Quantum Information & Computation*, 14(5&6):0361–0383, 2014.

[8] Bruno Grenet, Pascal Koiran, and Natacha Portier. On the complexity of the multivariate resultant. *Journal of Complexity*, 29(2):142 – 157, 2013.

[9] Julia Kempe and Oded Regev. 3-local Hamiltonian is QMA-complete. *Quantum Information & Computation*, 3(3):258–264, 2003.

[10] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, Boston, MA, USA, 2002.

[11] Klaus-Jörn Lange, Pierre McKenzie, and Alain Tapp. Reversible space equals deterministic space. *J. Comput. Syst. Sci.*, 60(2):354–367, April 2000.

[12] Chris Marriott and John Watrous. Quantum Arthur-Merlin games. *Computational Complexity*, 14(2):122–152, 2005.

[13] Daniel Nagaj, Pawel Wocjan, and Yong Zhang. Fast amplification of QMA. *Quantum Information & Computation*, 9(11):1053–1068, 2011.

[14] Norbert Schuch, Michael M. Wolf, Frank Verstraete, and J. Ignacio Cirac. Computational complexity of projected entangled pair states. *Phys. Rev. Lett.*, 98:140506, 2007.

[15] F. Verstraete and J. I. Cirac. Renormalization algorithms for quantum-many body systems in two and higher dimensions. 2004.

[16] John Watrous. Space-bounded quantum complexity. *J. Comput. Syst. Sci.*, 59(2):281–326, 1999.

[17] John Watrous. On the complexity of simulating space-bounded quantum computations. *Computational Complexity*, 12(1):48–84, 2003.

A Proof sketch of $\mathsf{PQP}^{O_{PEPS}}_{\|,\mathbf{classical}} = \mathsf{PP}$

Since $\mathsf{PP} \subseteq \mathsf{BQP}^{\mathsf{O_{PEPS}}}_{\|,\mathrm{classical}} \subseteq \mathsf{PQP}^{\mathsf{O_{PEPS}}}_{\|,\mathrm{classical}}$ [14], we only need to show that $\mathsf{PQP}^{O_{PEPS}}_{\|,\mathrm{classical}} \subseteq \mathsf{PP}$. In [14] it was noted that all PEPS can be seen as the output of a quantum circuit followed by a postselected measurement. Therefore $\mathsf{PQP}^{O_{PEPS}}_{\|,\mathrm{classical}}$ corresponds to the problems that can be decided by a quantum circuit, followed by a postselected measurement (since the queries to O_{PEPS} are classical and nonadaptive, we can compose them into one single postselection), followed by a measurement. In the YES case the measurement outputs 1 with probability at least c, whereas in the NO case the measurement outputs 1 with probability at most s, with $c > s$. The standard counting argument placing BQP inside PP then applies to this case as well; see for instance [1, Propositions 2 and 3].